

Fast Computation of Smith Forms of Sparse Matrices Over Local Rings *

Mustafa Elsheikh, Mark Giesbrecht, Andy Novocin
Cheriton School of Computer Science
University of Waterloo, Canada

B. David Saunders
Department of Computer and Information Sciences
University of Delaware, USA

*B. D. Saunders is supported by National Science Foundation Grants CCF-0830130, CCF-1018063. M. Elsheikh, M. Giesbrecht, and A. Novocin are supported by the Natural Sciences and Engineering Research Council of Canada, and MITACS (Canada).

Abstract

We present algorithms to compute the Smith Normal Form of matrices over two families of local rings. The algorithms use the *black-box* model which is suitable for sparse and structured matrices. The algorithms depend on a number of tools, such as matrix rank computation over finite fields, for which the best-known time- and memory-efficient algorithms are probabilistic.

For an $n \times n$ matrix A over the ring $F[z]/(f^e)$, where f^e is a power of an irreducible polynomial $f \in F[z]$ of degree d , our algorithm requires $\mathcal{O}(\eta de^2 n)$ operations in F , where our black-box is assumed to require $\mathcal{O}(\eta)$ operations in F to compute a matrix-vector product by a vector over $F[z]/(f^e)$ (and η is assumed greater than nde). The algorithm only requires additional storage for $\mathcal{O}(nde)$ elements of F . In particular, if $\eta = \tilde{\mathcal{O}}(nde)$, then our algorithm requires only $\tilde{\mathcal{O}}(n^2 d^2 e^3)$ operations in F , which is an improvement on known dense methods for small d and e .

For the ring $\mathbb{Z}/p^e\mathbb{Z}$, where p is a prime, we give an algorithm which is time- and memory-efficient when the number of nontrivial invariant factors is small. We describe a method for dimension reduction while preserving the invariant factors. The time complexity is essentially linear in $\mu nr e \log p$, where μ is the number of operations in $\mathbb{Z}/p\mathbb{Z}$ to evaluate the black-box (assumed greater than n) and r is the total number of non-zero invariant factors. To avoid the practical cost of conditioning, we give a Monte Carlo certificate, which at low cost, provides either a high probability of success or a proof of failure. The quest for a time- and memory-efficient solution without restrictions on the number of nontrivial invariant factors remains open. We offer a conjecture which may contribute toward that end.

Category: G.4. Mathematical Software Algorithm Design and Analysis

Category: I.1.4. Symbolic and Algebraic Manipulation Applications

Terms: Algorithms, Complexity, Performance.

Keywords: Local Principal Ideal Ring, Sparse Matrix, Polynomial Matrix, Integer Matrix, Smith Form, Black Box.

1 Introduction

We consider the problem of computing the Smith Normal Form (SNF) of sparse matrices over (commutative) local principal ideal rings (PIRs). The Smith form is a diagonalization of matrices which has many applications in diophantine analysis (Chou and Collins, 1982), integer programming (Hu, 1969), combinatorics (Wallis et al., 1972), determining the structure of Abelian groups (Newman, 1972) and class groups (Hafner and McCurley, 1989), computing Simplicial Homology (Dumas et al., 2003), in system theory (Kailath, 1980; McMillan, 1952), and in the study of symplectic spaces (Chandler et al., 2010).

The original work of Smith (1861) proved existence and uniqueness of the SNF for integer matrices. The generalization to PIRs is due to Kaplansky (1949).

The problem of computing the Smith form of a sparse matrix over a principal ideal ring presents several challenges. One approach is to simply compute the SNF over the global ring (i.e., $F[z]$ or \mathbb{Z}) and then reduce the result modulo the power of the prime ideal. The algorithm of Giesbrecht (2001) for SNF of a sparse matrix over \mathbb{Z} could be used, but the ultimate time requirement is essentially cubic (although space requirements are lower). An asymptotically faster algorithm along similar lines, but which requires considerably more space, is presented in (Eberly et al., 2007). The best known algorithm for dense matrices over $F[z]$ by (Storjohann, 2000, Prop 7.16) requires time essentially equal to matrix multiplication. However, it is not sensitive to sparsity.

On the other hand, computations over $F[z]$ and \mathbb{Z} suffer from coefficient growth which is not clearly necessary in a PIR. For example, over $\mathbb{Z}/p^2\mathbb{Z}$ where p is a prime, one might hope to perform all computations modulo p^2 , and not with integers larger than p^2 . Storjohann (2003) provides a fast algorithm using elimination in a PIR, but it is not sensitive to sparsity and requires time proportional to matrix multiplication. Wilkening and Yu (2011) demonstrates an algorithm for dense polynomial matrices over local rings, but offers no complexity analysis. Dumas et al. (2001) give black-box algorithms over \mathbb{Z} and locally at a prime, which however do not have a benefit when only a few invariant factors are nontrivial.

When dealing with sparse matrices we would like to preserve the sparsity of the input matrix, and introduce no fill-in. Thus we pursue *black-box* algorithms in the sense that the input matrix is only used for matrix-vector

products. The complexity of black-box algorithms is thus expressed in terms of the number of matrix-vector products used. Space requirements are kept to the storage of a few vectors. There has been great success in applying black-box methods over finite and arbitrary fields, starting with [Wiedemann \(1986\)](#), where the cost of many linear algebra problems has been reduced to computing a linear number of matrix-vector products. Our ultimate goal is then to add local Smith form to that list.

Specifically we will consider the local Artinian principal rings (also known as *special principal rings*) $\mathbb{Z}/p^e\mathbb{Z}$ for a prime p and positive exponent e , and $\mathbb{F}[z]/f^e\mathbb{F}[z]$, for irreducible $f \in \mathbb{F}[z]$. Let \mathbb{L} be a local Artinian principal ideal ring with a maximal prime ideal $p\mathbb{L}$. For any matrix $A \in \mathbb{L}^{n \times n}$, there exist unimodular matrices $U, V \in \mathbb{L}^{n \times n}$ and a diagonal matrix $S \in \mathbb{L}^{n \times n}$ such that $A = USV$, where

$$S = \text{diag}(\underbrace{1, \dots, 1}_{r_0}, \underbrace{p, \dots, p}_{r_1}, \dots, \underbrace{p^{e-1}, \dots, p^{e-1}}_{r_{e-1}}, 0, \dots, 0) \quad (1)$$

Definition 1. S is called the Smith form of A , and the diagonal elements are called the invariant factors of A .

Our goal throughout this paper, is to compute the multiplicities of the Smith form invariants, i.e., $\{r_0, r_1, \dots, r_{e-1}\}$ for given black-box matrices, and in particular, sparse matrices.

Our Contribution. For matrices over $\mathbb{F}[z]/(f^e)$, we present an algorithm which relies on computing ranks of related black-box matrices over \mathbb{F} , and give a complete complexity analysis. The key idea of our algorithm is a linear representation of polynomials in the ring $\mathbb{F}[z]/(f^e)$ as matrices over \mathbb{F} , and using rank computations over \mathbb{F} to discover the multiplicities of the Smith invariants. This reduction allows us to take advantage of the well-studied efficient algorithms for computing ranks of sparse matrices over fields rather than rings. The cost of our algorithm is $\mathcal{O}(\eta de^2 n)$ operations in \mathbb{F} , where each black-box evaluation costs η operations. Our approach takes a path similar to the linearization of [Kaltofen et al. \(1990\)](#) for matrices over $\mathbb{F}[z]$. This approach is also explored for dense matrices over local rings by [Wilkening and Yu \(2011\)](#). [Dumas et al. \(2009\)](#) used rank computations to discover multiplicities of characteristic polynomial factors for black-box matrices over fields.

The linearization idea, however, would not be applicable over the integers since there are no appropriate linear representations from $\mathbb{Z}/p^e\mathbb{Z}$ to $\mathbb{Z}^{n \times n}$.

Hence, it is necessary to develop different methods for the integer case. A useful approach to Smith form computation is to begin by determining which primes occur in the invariant factors and then compute the form locally at those primes. This has been done in several recent algorithms (Dumas et al., 2003; Saunders and Wan, 2004). A fully memory efficient, black-box algorithm for sparse and structured matrices has not been given, however, for lack of an efficient black-box algorithm for the Smith form locally at a prime.

Toward that end we give a black-box algorithm over $\mathbb{Z}/p^e\mathbb{Z}$, whose cost is essentially dominated by $\mu n e k$ for an $n \times n$ sparse matrix with μ nonzero entries, e being the largest exponent of the prime in the Smith form, and $k = \sum_{i=1}^{e-1} r_i$ being the number of nontrivial invariant factors (r_0 is *not* included). It is to be expected that the cost depends on both μn and e . The dependence on k , although unfortunate, is not completely unlikely. It is natural that it is easier to find Smith form for matrices with fewer number of non-trivial factors. In addition, both e and k are small in many cases of interest. Some applications with this property are discussed at the beginning of Section 3. The key idea of this algorithm is to apply a reduction in dimension to dispose the zero invariant factors, compute a nullspace basis of the reduced matrix to dispose the ones, and then determine the nontrivial invariant factors by dense elimination methods. This idea is applicable to the polynomial case as well. However, it is not interesting, since the complexity of this method has an extra factor of k over the rank-based method presented in Section 2.

Notation. Throughout this paper F denotes a field and L denotes a local ring. We use \mathbb{Z}_p to denote $\mathbb{Z}/p\mathbb{Z}$ and \mathbb{Z}_{p^e} to denote $\mathbb{Z}/p^e\mathbb{Z}$. In the complexity analysis, we count the *algebraic complexity* in the base field, i.e. we assume that all operations in the base field have a unit cost. We use “soft-Oh” to hide the logarithmic factors. We say that $f \in \mathcal{O}^\sim(g)$ if there exists a constant c such that $f \in \mathcal{O}(g \log^c g)$. We use $M(n)$ to denote the number of operations in the base field required to multiply two polynomials of degree at most n . Finally, we use $\mathcal{O}(n^\omega)$ to denote the matrix multiplication exponent, e.g., $\omega \leq 2.372$ using Coppersmith and Winograd (1990).

Roadmap. The rest of this paper is organized as follows. Section 2 contains the algorithm and analysis for Smith form over $F[z]/(f^e)$. In Section 3, our *nullspace* algorithm for Smith form over \mathbb{Z}_{p^e} is given after some discussion of applications, a development of preconditioners for the problem. A Monte Carlo verification method is given that can be of use with small primes. A conjecture is also given, that may shed some light on the problem when there

are many nontrivial invariants. Finally, Section 4 is a brief summary.

2 Smith form over $\mathbb{F}[z]/(f^e)$

In this section we present an algorithm to compute the local Smith form of a sparse polynomial matrix. Throughout this section, let \mathbb{F} be a field, $\mathbb{L} = \mathbb{F}[z]/(f^e)$ where $f \in \mathbb{F}[z]$ is irreducible of degree d , and $e \in \mathbb{Z}_{>1}$. The ideals in this ring are of the form $f^i\mathbb{L}$ for $0 \leq i < e$ and the RHS of equation (1) becomes

$$\text{diag}(\underbrace{1, \dots, 1}_{r_0}, \underbrace{f, \dots, f}_{r_1}, \dots, \underbrace{f^{e-1}, \dots, f^{e-1}}_{r_{e-1}}, 0, \dots, 0) \quad (2)$$

Our goal is efficiently compute the multiplicities: $\{r_0, r_1, \dots, r_{e-1}\}$. The approach is to embed the ring $\mathbb{L}^{n \times n}$ in the ring $\mathbb{F}^{nde \times nde}$ and reduce the computation to finding ranks of matrices in the base field, \mathbb{F} , where known fast black-box algorithms can be used.

2.1 Embedding of $\mathbb{L}^{n \times n}$ in $\mathbb{F}^{nde \times nde}$

In this section, we describe the classical embedding of \mathbb{L} into $\mathbb{F}^{de \times de}$, and how properties of matrices over \mathbb{L} are revealed by their images over \mathbb{F} . First, define the map $\varphi_e : \mathbb{L} \rightarrow \mathbb{F}^{de \times de}$ as follows. Suppose $f^e = a_0 + a_1z + \dots + a_{de-1}z^{de-1} + z^{de}$, with a companion matrix

$$C_{f^e} = \begin{pmatrix} 0 & 0 & \cdots & -a_0 \\ 1 & \ddots & & -a_1 \\ \vdots & \ddots & \ddots & \vdots \\ 0 & \cdots & 1 & -a_{de-1} \end{pmatrix}.$$

Define $\varphi_e(z) = C_{f^e}$, and $\varphi_e(z^i) = \varphi_e(z)^i$. By linearity, extend φ_e to all elements of $g = g_0 + g_1z + \dots + g_{de-1}z^{de-1} \in \mathbb{L}$ such that $\varphi_e(g) \in \mathbb{F}^{de \times de}$:

$$\varphi_e(g) = g(C_{f^e}) = g_0I + g_1C_{f^e} + g_2C_{f^e}^2 + \dots + g_{de-1}C_{f^e}^{de-1}.$$

It is straightforward to verify that φ_e is a ring isomorphism between \mathbb{L} and $\mathbb{F}[C_{f^e}]$.

Lemma 2. $\text{rank}(\varphi_e(f^i)) = d(e - i)$ for $0 \leq i \leq e$.

Proof. Since $f(C_{f^e})$ acts as multiplication by $f \bmod f^e$, it has null vectors which are images of polynomials in $f^{e-1}\mathbf{L}$. This is a vector space of dimension d , and hence $\text{rank}(f(C_{f^e})) = de - d$. Also, $f(C_{f^e})$ has minimal polynomial x^e , whence

$$\varphi_e(f) \sim N_f = \begin{pmatrix} 0_d & I_d & \cdots & 0_d \\ \vdots & \ddots & & \vdots \\ \vdots & \ddots & \ddots & I_d \\ 0_d & \cdots & \cdots & 0_d \end{pmatrix},$$

where $I_d, 0_d \in \mathbb{F}^{d \times d}$ are identity and zero matrices respectively. The rank $\varphi_e(f^i) = \varphi_e(f)^i$ is now evident from the structure of the nilpotent N_f . \square

We extend the map φ_e to $n \times n$ matrices over \mathbf{L} . For every $A \in \mathbf{L}^{n \times n}$, $\varphi_e(A)$ is a $nde \times nde$ matrix over \mathbb{F} , where every entry $a_{i,j}$ of A is replaced by the $de \times de$ block $\varphi_e(a_{i,j})$. Applying φ_e to (2), we get

$$\begin{aligned} \varphi_e(S) = & \text{diag}(\underbrace{\varphi_e(1), \dots, \varphi_e(1)}_{r_0}, \underbrace{\varphi_e(f), \dots, \varphi_e(f)}_{r_1}, \dots, \\ & \underbrace{\varphi_e(f^{e-1}), \dots, \varphi_e(f^{e-1})}_{r_{e-1}}, 0, \dots, 0) \in \mathbb{F}^{nde \times nde}. \end{aligned} \quad (3)$$

Lemma 3. *If $U \in \mathbf{L}^{n \times n}$ is invertible, then $\varphi_e(U) \in \mathbb{F}^{nde \times nde}$ is invertible.*

Proof. If U is invertible, then there exists a $W \in \mathbf{L}^{n \times n}$ such that $UW = I$, and $\varphi_e(U)\varphi_e(W) = \varphi_e(I)$. But $\varphi_e(I) = I_{nde}$, so $\varphi_e(U)$ has inverse $\varphi_e(W)$. \square

We now establish the property relating multiplicities in the invariant factors of $A \in \mathbf{L}^{n \times n}$ to the rank of $\varphi_e(\mathbf{L})$.

Theorem 4. *Let $A \in \mathbf{L}^{n \times n}$ have Smith form in (2), then $\text{rank}(\varphi_e(A)) = der_0 + d(e-1)r_1 + \cdots + dr_{e-1}$.*

Proof. There exist unimodular matrices $U, V \in \mathbf{L}^{n \times n}$ such that $UAV = S$. By isomorphism, $\varphi_e(U)\varphi_e(A)\varphi_e(V) = \varphi_e(S)$. By Lemma 3, $\varphi_e(U), \varphi_e(V)$ are invertible and thus $\text{rank}(\varphi_e(A)) = \text{rank}(\varphi_e(S))$. In (3), $\varphi_e(S)$ is a block diagonal matrix, so

$$\text{rank}(\varphi_e(S)) = \sum_{i=0}^{e-1} \text{rank}(\varphi_e(f^i)).$$

The proof then follows from Lemma 2. \square

A consequence of Theorem 4 is that, for $1 \leq \ell \leq e$, we have $\text{rank}(\varphi_\ell(A \bmod f^\ell)) = d\ell r_0 + d(\ell-1)r_1 + \dots + dr_{\ell-1}$. For example, $\text{rank}(\varphi_1(A \bmod f)) = dr_0$. In general, we have the following corollary. The proof is left to the reader.

Corollary 5. *Let $\rho_{\ell-1}$ denote $\text{rank}(\varphi_e(A \bmod f^\ell))$, $1 \leq \ell \leq e$. Then*

$$\begin{pmatrix} d & 0 & \cdots & 0 \\ 2d & d & \cdots & 0 \\ \vdots & \ddots & \ddots & \vdots \\ ed & \cdots & 2d & d \end{pmatrix} \begin{pmatrix} r_0 \\ r_1 \\ \vdots \\ r_{e-1} \end{pmatrix} = \begin{pmatrix} \rho_0 \\ \rho_1 \\ \vdots \\ \rho_{e-1} \end{pmatrix}. \quad (4)$$

This system may be solved in linear time. Let $\sigma_1 = \rho_1$ and $\sigma_i = \rho_i - \rho_{i-1}$, for $1 < i < e$. Then the σ_i are the prefix sums of the r_i , so $r_1 = \sigma_1$ and $r_i = \sigma_i - \sigma_{i-1}$, for $1 < i < e$.

Next we consider how to efficiently compute $\{\rho_0, \rho_1, \dots, \rho_{e-1}\}$ for a given black-box matrix.

2.2 Black-box for the embedding

Given a black-box for $A \in \mathbb{L}^{n \times n}$, over $\mathbb{L} = \mathbb{F}[z]/(f^e)$, we can easily construct a black-box for $\varphi_\ell(A \bmod f^\ell)$, for all $\ell \leq e$, at not much higher cost.

First, we define the black-box model cost model, and then show how to perform black-box computations under φ_e transformations efficiently.

Definition 6. *Let $A \in \mathbb{L}^{n \times n}$ be a sparse matrix over $\mathbb{L} = \mathbb{F}[z]/(f^e)$, where $f \in \mathbb{F}[z]$ is monic and irreducible of degree d . The black-box for A is a mapping $\mathbb{L}^n \rightarrow \mathbb{L}^n$ such that for all $v \in \mathbb{L}^n$, $Av \in \mathbb{L}^n$ can be computed with η operations in \mathbb{F} . We assume throughout that $\eta \geq nde$.*

Lemma 7. *Suppose we are given a black-box for $A \in \mathbb{L}^{n \times n}$, where $\mathbb{L} = \mathbb{F}[z]/(f^e)$ as above. Let $\ell \in \{1, \dots, e\}$ and $\hat{v} \in \mathbb{F}^{d\ell n}$ with unique pre-image $v \in \mathbb{F}[z]/(f^\ell)$. Then we can compute $\varphi_\ell(A \bmod f^\ell)\hat{v} \in \mathbb{F}^{d\ell n}$ with $\mathcal{O}(\eta + nM(de))$ operations in \mathbb{F} .*

Proof. Assume that $\hat{v} \in \mathbb{F}^{d\ell n}$ is labelled as:

$$\hat{v} = (\hat{v}_{1,0}, \dots, \hat{v}_{1,d\ell-1}, \hat{v}_{2,0}, \dots, \hat{v}_{2,d\ell-1}, \dots, \hat{v}_{n,0}, \dots, \hat{v}_{n,d\ell-1}).$$

Construct the vector $v = (v_1, \dots, v_n) \in \mathbb{L}^n$, where $v_i = \sum_{0 \leq j < d\ell} \hat{v}_{i,j} z^j \in \mathbb{F}[z]$. Now, compute $w = Av \bmod f^\ell \in \mathbb{L}^n$ using η operations for the black-box

evaluation plus $\mathcal{O}(n \mathbf{M}(de))$ operations in \mathbb{F} . Let $w = (w_1, \dots, w_n)$. Assume $w_i = \sum_{0 \leq j < d\ell} \hat{w}_{i,j} z^j \in \mathbb{F}[z]$. Then

$$\hat{w} = (\hat{w}_{1,0}, \dots, \hat{w}_{1,d\ell-1}, \dots, \hat{w}_{n,0}, \dots, \hat{w}_{n,d\ell-1}).$$

□

Our algorithm for computing the Smith form of a matrix $A \in \mathbb{L}^{n \times n}$ given by a black-box is now straightforward. Using Theorem 4 and Lemma 7 we reduce the computation of ρ_i 's in (4) to computing ranks of matrices over the ground field \mathbb{F} , which can be accomplished using existing fast and memory-efficient black-box algorithms over fields, e.g. Wiedemann's algorithm.

Algorithms for computing the rank of a black-box matrix over a field are developed by Wiedemann (1986), and refined in subsequent work of Kaltofen and Saunders (1991), Eberly (2004), and others. If the input matrix is in $\mathbb{F}^{n \times n}$ and the black-box evaluation requires η operations in \mathbb{F} , then the rank algorithms require $\mathcal{O}(n\eta)$ operations in \mathbb{F} . They are probabilistic, and return the correct rank with controllably high probability on any input. We will assume that an appropriate choice of black-box rank method is made, and note that there is considerable difference in their effectiveness in practice and over various ground fields.

Algorithm 1. *Smith invariants in $\mathbb{L} = \mathbb{F}[z]/(f^e)$, where $f \in \mathbb{F}[z]$ is irreducible of degree d .*

Input: Black-box for $A \in \mathbb{L}^{n \times n}$.

Output: r_0, \dots, r_{e-1} such that r_i is the multiplicity of f^i in the Smith Form of A , and the multiplicity of 0 is $n - \sum_i r_i$.

1. For all $\ell \in \{1, \dots, e\}$, invoke a black-box rank algorithm on the black-box for $\varphi_\ell(A \bmod f^\ell) : \mathbb{F}^{d\ell n} \rightarrow \mathbb{F}^{d\ell n}$. Let $\rho_{\ell-1} = \text{rank}(\varphi_\ell(A \bmod f^\ell))$.
2. Solve (4) for r_0, \dots, r_{e-1} .
3. Return r_0, \dots, r_{e-1} .

Theorem 8. *Algorithm 1 is correct, and requires $\mathcal{O}(\eta de^2 n)$ operations in \mathbb{F} . The space requirement of the algorithm is $\mathcal{O}(\text{den})$ elements in \mathbb{F} .*

Proof. The correctness follows from the results and discussion in this section. We analyze the time and space complexity of step (1), which dominates. This step requires $\mathcal{O}(de^2 n)$ black-box evaluations, and storage for $\mathcal{O}(nde)$ elements in \mathbb{F} . □

3 Smith form over \mathbb{Z}_{p^e}

In our experience, most Smith normal forms of integer matrices in practice involve relatively few non-trivial factors, i.e., most of the invariant factors are 1's or 0's. The algorithm of this section addresses that situation.

For example, in some work on computing homology of simplicial complexes, [Dumas et al. \(2000, 2001\)](#); [Babson et al. \(1999\)](#); [Björner and Welker \(1999\)](#), large boundary matrices arose. One of the most challenging Smith forms to compute at the time was a 135135 by 270270 matrix which turned out to have 133991 ones, 220 3's, and 924 zeroes as the invariants. Other examples in that study were also large but with even fewer nontrivial invariants. Most often in homology computation, the number of 1's (the Betti number) greatly exceeds the number of non-trivial entries, it seems.

For another example, recently in the study of symplectic 3 spaces, Smith form computations have been desired of some rather large matrices [Chandler et al. \(2010\)](#); [Chandler \(2011\)](#). For these matrices it is conjectured that only one prime will appear in the invariant factors (other than the largest) and indeed only the local Smith form at that prime is of interest. Furthermore, the conjectured structure predicts only a few nontrivial invariant factors. We denote these examples as W3- q , where q is a prime power and the Smith form modulo q is desired. W3- q is a $\{0, 1\}$ -matrix of size approximately $q^3 \times q^3$ with about q^4 nonzero entries. The current challenge is to compute Smith form of W3-64 and W3-81. The algorithm described here is designed to handle this case.

For a prime p and an exponent $e \in \mathbb{Z}_{>1}$, let φ be the natural projection $\mathbb{Z}_{p^e} \rightarrow \mathbb{Z}_p$, which extends naturally to $\varphi : \mathbb{Z}_{p^e}^{n \times n} \rightarrow \mathbb{Z}_p^{n \times n}$ by element-wise mapping. Note that $x \in \mathbb{Z}_{p^e}$ is a unit if and only if $\varphi(x) \neq 0$. Likewise, $A \in \mathbb{Z}_{p^e}^{n \times n}$ is unimodular if and only if $\varphi(A)$ is unimodular.

3.1 Nullspace method

Let us introduce the approach by way of a sketched example. Suppose A is a matrix over \mathbb{Z}_{p^5} . Further suppose A is 100 by 100 and

$$A \sim \text{diag}(1, 1, \dots, 1, p, p, p, p^3, p^4, 0, 0, \dots, 0),$$

with 45 ones and 50 zeroes. This approximates on a small scale the pattern of invariants we expect to see on W3- q . First, a reduction in dimension allows

us to reduce A to an $\ell \times \ell$ matrix $\rho(A)$ having the same nonzero invariants, where ℓ is the (max) rank, or slightly larger. We illustrate with $\ell = 52$:

$$\rho(A) \sim S = \text{diag}(1, 1, \dots, 1, p, p, p, p^3, p^4, 0, 0).$$

Over \mathbb{Z}_p , the nullspace basis N' of $\varphi(\rho(A))$ (N' has 7 columns) is equivalent to that of S . Let E' be the last 7 columns of the 52×52 identity matrix. Let E and N be arbitrary embeddings of E' and N' in $\mathbb{Z}_{p^5}^{52 \times 7}$, such that $\varphi(E) = E', \varphi(N) = N'$. Thus $\rho(A)N$ and SE are multiples of p and

$$\rho(A)N \sim SE = \text{diag}(p, p, p, p^3, p^4, 0, 0).$$

In summary, the algorithm is to apply a reduction in dimension to dispose of zeroes, compute nullspace basis N to dispose of ones, and determine the nontrivial invariants by computing Smith form of AN using dense methods. AN is an $n \times k$ matrix, where k is the number of nontrivial invariants (or slightly larger – ℓ - rank mod 2).

Reduction in dimension is a frequent tool and has been used for Smith form, for example, in [Dumas et al. \(2001\)](#). But then their computation proceeds without disposing of the unit invariant factors. Thus the time complexities below, otherwise similar to theirs, differ in that we replace a rank factor ℓ by the number of nontrivial invariants, k .

3.2 Probabilistic dimension reduction

Let $A \in \mathbb{Z}_{p^e}^{n \times n}$, for which we have a fast black-box. Let A have a Smith form $\text{diag}(s_1, \dots, s_r, 0, \dots, 0) \in \mathbb{Z}_{p^e}^{n \times n}$. Our goal in this subsection is to construct $\rho(A)$. That is, given such an A and a $\ell \in \{1, \dots, n\}$, to construct a black-box of similar cost for a matrix $B \in \mathbb{Z}_{p^e}^{\ell \times \ell}$ which has Smith form $\text{diag}(s_1, \dots, s_\ell)$, i.e., with the initial invariant factors of A .

Notationally, for integers n and $k \leq n$, let \mathcal{C}_k^n be the set of k -tuples of distinct elements (in increasing order) of $\{1, \dots, n\}$. For a matrix $B \in \mathbb{L}^{n \times n}$, and $\sigma, \tau \in \mathcal{C}_k^n$, define $B_{(\sigma, \tau)}^{(\sigma)}$ as the (σ, τ) minor of B , i.e., the determinant of the $k \times k$ submatrix of B with rows from σ and columns from τ . We use script letters, e.g. $\mathfrak{D}, \mathfrak{T}$, to denote matrices with indeterminate entries.

We use techniques similar to that derived in [Giesbrecht \(2001\)](#) with *scaled Toeplitz matrix conditioners*. For indeterminates $\Lambda = \{v_i, w_i, y_i\}$, let $\mathfrak{D}_1 =$

$\text{diag}(v_1, \dots, v_n)$, $\mathfrak{D}_2 = \text{diag}(w_1, \dots, w_n)$, and \mathfrak{T} be a generic Toeplitz matrix:

$$\mathfrak{T} = \begin{pmatrix} y_n & y_{n+1} & \cdots & y_1 \\ \vdots & y_n & \ddots & \vdots \\ y_{n-2} & & \ddots & y_{n+1} \\ y_{2n-1} & y_{2n-2} & \cdots & y_n \end{pmatrix} \quad (5)$$

Lemma 9. *Let $\mathfrak{B} = \mathfrak{D}_1 \mathfrak{T} \mathfrak{D}_2$ in the indeterminates Λ , as in (5). Let $k \in \{1, \dots, n\}$ and $\sigma = (\sigma_1, \dots, \sigma_k)$, $\tau = (\tau_1, \dots, \tau_k) \in \mathcal{C}_k^n$.*

- (i) $\mathfrak{T}(\sigma) \in \mathbb{Z}[\Lambda]$ has content 1;
- (ii) $\mathfrak{B}(\sigma) = v_{\sigma_1} \cdots v_{\sigma_k} w_{\tau_1} \cdots w_{\tau_k} \mathfrak{T}(\sigma)$.

Proof. Part (i) is from (Giesbrecht, 2001, Lemma 1.3) and part (ii) follows easily from the Cauchy-Binet formula. \square

Note that $\mathfrak{B}(\sigma)$ uniquely identifies which minor of \mathfrak{B} was selected.

Lemma 10. *Let $A \in \mathbb{Z}^{n \times n}$, and $\mathfrak{B}_1, \mathfrak{B}_2$ be $n \times n$ matrices of distinct indeterminates from a set Λ , of the form (5). Then $\mathfrak{A} = \mathfrak{B}_1 A \mathfrak{B}_2$ is such that for $1 \leq k \leq n$, the content of $\psi_k = \mathfrak{A}(\sigma) \in \mathbb{Z}[\Lambda]$ equals Δ_k , the k th determinantal divisor of A .*

Proof. By the Cauchy-Binet formula we have

$$\mathfrak{A}(\sigma) = \sum_{\tau \in \mathcal{C}_k^n} \mathfrak{B}_1(\sigma) A(\tau) \mathfrak{B}_2(\tau).$$

Thus $\mathfrak{A}(\sigma)$ is a sum of polynomials of content 1, with distinct indeterminates, one for each $k \times k$ minor of A , times the value of that minor. Hence it must have content equal to the GCD of all $k \times k$ minors of A , which is equal to the k th determinantal divisor. \square

Theorem 11. *Let $A \in \mathbb{Z}^{n \times n}$, $p \geq 6n^2\xi$ a prime, and $\xi \geq 2$. Let $B_1, B_2 \in \mathbb{Z}^{n \times n}$ be formed by a random assignment of variables in $\mathfrak{B}_1, \mathfrak{B}_2$ in (5) respectively, where choices are made uniformly from $L = \{0, \dots, 6n^2\xi - 1\}$, and $\hat{A} = B_1 A B_2$. Then with probability at least $1 - 1/\xi$, for all $1 \leq k \leq n$, the order of p in Δ_k , the k th determinantal divisor of A equals the order of p in $\hat{A}(\sigma)$.*

Proof. Let ψ_k be as in Lemma 10, which has content equal to the k th determinantal divisor Δ_k of A . Observe from our construction that $\deg \psi_k \leq 6k \leq 6n$ (the total degree of a $k \times k$ minor of an indeterminate Toeplitz is $\leq k$). Thus, with values selected as described, by the Schwartz-Zippel Lemma (Zippel, 1979; Schwartz, 1980), we have that ψ_k/Δ_k is a polynomial in the entries of matrices B_1, B_2 and

$$\text{prob} \{(\psi_k/\Delta_k) \not\equiv 0 \pmod{p}\} \geq 1 - \frac{6n}{6n^2\xi}.$$

This is exactly the probability that the order of p in Δ_k equals the order of p in the leading $k \times k$ minor of \hat{A} . The probability that this happens for all k , from $1 \leq k \leq n$ is at least $(1 - 1/(n\xi))^n \geq 1 - 1/\xi$. \square

Corollary 12. *Let $p \geq 6n^2\xi$ be prime, for a $\xi > 1$, and $e \geq 1$, and suppose $A \in \mathbb{Z}_{p^e}^{n \times n}$ has (local) Smith form $\text{diag}(s_1, \dots, s_n) \in \mathbb{Z}_{p^e}^{n \times n}$. Let $B_1, B_2 \in \mathbb{Z}_{p^e}^{n \times n}$ be formed by a random assignments of variables in $\mathfrak{B}_1, \mathfrak{B}_2 \in \mathbb{Z}^{n \times n}$ in (5) respectively, where choices are made uniformly from $L = \{0, \dots, 6n^2\xi - 1\} \pmod{p^e}$. Let $\hat{A} = B_1AB_2 \in \mathbb{Z}_{p^e}^{n \times n}$, and for $1 \leq k \leq n$ let \hat{A}_k be the leading $k \times k$ submatrix of \hat{A} . Then with probability at least $1 - 1/\xi$, for all $k \in \{1, \dots, n\}$, the local Smith form of \hat{A}_k is $\text{diag}(s_1, \dots, s_k) \in \mathbb{Z}_{p^e}^{k \times k}$.*

Proof. The Smith form of A equals the Smith form of any $\tilde{A} \in \mathbb{Z}^{n \times n}$ with $\tilde{A} \equiv A \pmod{p^e}$, reduced modulo p^e (the only non-units will be powers of p after the reduction). Thus, Theorem 11 implies that the order of p in the k th determinantal divisor of A equals the order of p in the leading $k \times k$ minor of \tilde{A} , for all k , with probability at least $1 - 1/\xi$. This implies that $\hat{A}_k = \tilde{A}_k \pmod{p^e}$ will have Smith form (s_1, \dots, s_k) for all $1 \leq k \leq n$ where \tilde{A}_k is the leading $k \times k$ minor of \tilde{A} , since $\Delta_k = s_1 \cdots s_k$ for $1 \leq k \leq n$. \square

Computationally, if we know that rank of A is at most m , then we can work with truncated random scaled Toeplitz matrices $B_1 \in \mathbb{Z}_{p^e}^{m \times n}$ and $B_2 \in \mathbb{Z}^{n \times m}$. Then Corollary 12 implies that $\hat{A} = B_1AB_2 \in \mathbb{Z}_{p^e}^{m \times m}$ has the same non-zero invariant factors as A .

Working with Small Primes. The conditions for Corollary 12 require $p \geq 6n^2\xi$. For smaller primes the algorithm may well work, but appears much more difficult to prove. The following method may be used to remedy this.

The approach is to replace \mathbb{Z} by a subring of the ring of algebraic integers in a number field of degree $\eta = \lceil \log_p(6n^2\xi) \rceil$ over \mathbb{Q} , in which p is inert. Specifically, let $\Gamma \in \mathbb{Z}[y]$ have degree at least η be such that $\Gamma \bmod p$ is irreducible in $\mathbb{Z}_p[y]$. Let $\gamma \in \mathbb{C}$ be a root of Γ . Then $\mathbb{Z}[\gamma]$ is such that $\mathbb{Z}[\gamma]/(p^e)$ is a local ring which contains \mathbb{Z}_{p^e} , and such that the residue class field $\mathbb{Z}[\gamma]/(p)$ contains more than $6n^2\xi$ elements. We call $\mathbb{Z}[\gamma]/(p^e)$ the *Galois ring* with p^η elements, and denote it by $\text{GR}(p^e, \eta)$ (see [McDonald \(1974\)](#)). Like \mathbb{Z}_{p^e} , $\text{GR}(p^e, \eta)$ is a local principal ideal ring with maximal prime ideal generated by p .

Analogues of Theorem 11 (over $\mathbb{Z}[\gamma]$) and Corollary 12 (over $\text{GR}(p^e, \eta)$) can be proven similarly. We state the latter formally, but leave the proofs to the reader.

Corollary 13. *Let p be prime, $e \geq 1$, $\xi \geq 1$ and $\eta = \lceil \log_p(6n^2\xi) \rceil$. Let $\text{GR}(p^e, \eta) = \mathbb{Z}[y]/(\Gamma)$ for $\Gamma \in \mathbb{Z}[y]$ of degree η which is irreducible modulo p . Suppose $A \in \mathbb{Z}_{p^e}^{n \times n}$ has (local) Smith form $\text{diag}(s_1, \dots, s_n) \in \mathbb{Z}_{p^e}^{n \times n}$. Let $B_1, B_2 \in \text{GR}(p, \eta)^{n \times n}$ be formed by random assignments of indeterminates in $\mathfrak{B}_1, \mathfrak{B}_2$ in (5) respectively, where choices are made uniformly from $L = \{\sum_{0 \leq i < \eta} \alpha_i y^i : \alpha_i \in \{0, \dots, p-1\}\} \bmod p^e$. Let $\hat{A} = B_1 A B_2 \in \text{GR}(p, \eta)^{n \times n}$, and for $1 \leq k \leq n$ let \hat{A}_k be the leading $k \times k$ submatrix of \hat{A} . Then with probability at least $1 - 1/\xi$, for all $k \in \{1, \dots, n\}$, the local Smith form of \hat{A}_k is $\text{diag}(s_1, \dots, s_k) \in \mathbb{Z}_{p^e}^{k \times k}$.*

3.3 Probabilistic validation of dimension reduction

This section might be titled, “Escaping the tyranny of Schwartz-Zippel.” It is frequently the case, as is exemplified in the previous section, that an argument for a favourable probability based on the Schwartz-Zippel Lemma requires an inconveniently large set for random assignments. Experience in practice demonstrates that far smaller sets suffice virtually always. In this section we give a method to have a provably low probability of failure while making no assumptions about the basic preconditioner.

For a matrix A let $s_k(A)$ denote the k -th invariant factor of A (in the order in which $s_i(A) \mid s_{i+1}(A)$, for $1 \leq i < n$). Let $S_k(A)$ denote the leading $k \times k$ submatrix of the Smith form of A , $S_k(A) = \text{diag}(s_1(A), s_2(A), \dots, s_k(A))$. Let $[A, B]$ denote the side by side join of two conformable matrices.

Theorem 14. *Let $A \in \mathbb{L}^{m \times n}$, $Q \in \mathbb{L}^{n \times k}$, and $y \in \mathbb{L}^n$, with $k \leq \min(m, n)$ and the entries of y being uniform random variables over \mathbb{L} . Then AQ has*

the first k invariant factors of A with high probability if AQ and $A[Q, y]$ have the same first k invariant factors. To be precise, we have the following conditional probability:

$$\text{prob}(S_k(A) = S_k(AQ) \mid S_k(AQ) = S_k([AQ, Ay])) \geq 1 - 1/p.$$

Proof. Note that $S_k(A) = S_k(B)$ if and only if the i -th determinantal divisors are equal: $D_i(A) = D_i(B)$, $i \in 1, \dots, k$, where $D_i(A)$ denotes the GCD of all $A_{\tau}^{(\sigma)}$, $\sigma, \tau \in \mathcal{C}_i^n$. We will call any $i \times i$ minor equal to $D_i(A)$ (up to unit multiple) a *witness* for the determinantal divisor $D_i(A)$. Note that the GCD of a set of elements of \mathbb{L} is a member of the set (up to unit multiple), so every determinantal divisor has at least one witness. For example, for a Smith form S , the j -th determinantal divisor, if not a unit, has exactly one witness, $D_j(S) = S_{\iota(j)}^{(\iota(j))}$, where we define ι by $\iota(j) = (1, 2, \dots, j)$.

Without loss of generality, suppose for notational simplicity that $k \leq n \leq m$. Let $A = USV$ be the Smith form factorization of A with $S = \text{diag}(s_1, \dots, s_n)$ and U, V unimodular. We are concerned with the invariants of A, AQ , and $A[Q, y]$. Multiplication by unimodular U^{-1} does not affect invariants, so let us consider SV, SVQ , and $S[VQ, Vy]$. Because V is unimodular, Vy is a uniform random variable if y is. Also, we have made no conditions on Q , so we simplify notation, substituting Q for VQ and y for Vy . In other words, without loss of generality we may assume $A = S$ is in Smith form.

We will show the contrapositive of our proposition. Suppose j is the first index at which $S_j(A) \neq S_j(AQ)$. Then $A_{\iota(j)}^{(\iota(j))} \neq A_{\iota(j)}^{(\iota(j))} Q_{\sigma}^{(\iota(j))}$, for all $\sigma \in \mathcal{C}_j^k$. Otherwise such a minor would witness equality of $S_j(A)$ and $S_j(AQ)$. It follows that $p \mid Q_{\sigma}^{(\iota(j))}$.

Because j is the first such case, there must be an $(j-1) \times (j-1)$ minor $Q_{\tau}^{(\iota(j-1))}$ which is a unit (not divisible by p), where $\iota(j-1), \tau \in \mathcal{C}_{j-1}^k$. Let τ' denote $\tau \cup \{k+1\}$, a column index set for $[Q, y]$ which includes the last column. Then the expansion of the ι by τ' minor of $A[Q, y]$ has as coefficient of y_j the term $A_{\iota(j)}^{(\iota(j))} Q_{\tau}^{(\iota(j-1))}$, which is a unit. Thus, for each setting of y_1, \dots, y_{j-1} , there is at most one value modulo p or y_j for which $p \mid [Q, y]_{\tau'}^{(\iota(j))}$. For all other values, this minor witnesses that $D_j(AQ) \neq D_j(A[Q, y])$. Thus when $S_k(A) \neq S_k(AQ)$ there is at most a $1/p$ chance that $S_k(A[Q, y])$ agrees with $S_k(AQ)$. \square

The following corollary allows us to verify or disprove the success of preconditioners such as those used in dimension reduction. This works when the

theory of the preconditioner’s probability of success is invalidated because of an insufficiently large set from which random values are chosen. Thus in practice for a small prime, one can skip the domain extension described at the end of Section 3.2. Indeed, no theory of the preconditioner is required at all. One can try to “get lucky” and skip preconditioners altogether. For instance, apply the corollary where PAQ selects the leading $k \times k$ submatrix of A . If the method validates, then the Smith form was found economically, otherwise try a more thorough preconditioning.

Corollary 15 (projection verification). *Let matrix $A \in \mathbb{Z}^{n \times n}$ be given and heuristic preconditioners $P \in \mathbb{Z}^{k \times n}$ and $Q \in \mathbb{Z}^{n \times k}$. Choose $R_1 \in \mathbb{Z}^{n \times c}$, $R_2 \in \mathbb{Z}^{c \times n}$ at random. Let $B = \begin{pmatrix} PAQ & PAR_1 \\ R_2AQ & R_2AR_1 \end{pmatrix} \in \mathbb{Z}^{k+c \times k+c}$. If $S_k(PAQ) = S_k(B)$ then these are the first k invariant factors of A with probability greater than $1 - 2/p^c$.*

Proof. Let $C = \begin{pmatrix} PA \\ R_2A \end{pmatrix}$ so that $B = [CQ, CR_1]$. Note that the condition $S_k(PAQ) = S_k(B)$ implies that all minors of B containing PAQ have those invariants. In particular B and its left side CQ have the same invariants. So Theorem 14 applies c times for each of the columns of R_1 . Because these are independent random vectors the probability that $S_k(CQ) = S_k([CQ, CR_1])$ when $S_k(CQ) \neq S_k(C)$ is at most $1/p^c$. Then apply Theorem 14 to A, PA, C on the left c times for the c rows of R_2 . Again, the probability of an unfortunate equality is at most $1/p^c$, and otherwise $S_k(PA) = S_k(A)$ is validated. Thus the overall probability of success is at least $(1 - 1/p^c)^2 > 1 - 2/p^c$. \square

The idea to use some random dense rows in preconditioning is widespread. It was used already in (Wiedemann, 1986, proof of theorem 1), and in a sense it is the basis of block iterative methods. The idea to obtain a good probability (especially for small primes) by solving a sparse problem twice, first without the few random dense rows and/or columns then with them, was used in Saunders and Youse (2009). For integer lattices, there are also some similarities to the additive preconditioners of Eberly et al. (2000), and the lattice compression of Chen and Storjohann (2005), especially in the analysis for small primes dividing invariant factors.

The parameter c can be adjusted to get the desired degree of certainty. For example, $c = 21$ ensures probability of failure less than one in a million, since $2/p^c \leq 1/2^{20} \leq 10^{-6}$ in that case. Also, one can pad with more random

rows and columns to improve weak preconditioners. Thus if $c = 30$ is used and the matrix with the first 9 of those columns and rows has the same Smith form as the matrix with all 30 random rows and 30 random columns adjoined, we have computed the Smith form with error expected less than once in a million trials. In effect, we have corrected for some weakness in the heuristic preconditioners with 9 extra rows and columns and then verified with 21 more.

3.4 Algorithm for the Smith Normal Form

After reducing the dimension to a value at or near the number of nonzero invariant factors, the following algorithm is applied. Recall that φ is the natural projection $\mathbb{Z}_{p^e} \rightarrow \mathbb{Z}_p$.

Algorithm 2. *Smithpe-nullspace*

Input: a black-box for $B \in \mathbb{Z}_{p^e}^{n \times n}$, and a bound ℓ for the number of nonzero invariant factors

Output: S , the Smith form of B .

0. Set A to the dimension reduction of B to $\ell \times \ell$.
1. Let $r_0 = \text{rank}(\varphi(A))$ over \mathbb{Z}_p . The nullity of $\varphi(A)$ is then $k = \ell - r_0$.
2. Compute $N' \in \mathbb{Z}_{p^e}^{\ell \times k}$, a lifting to \mathbb{Z}_{p^e} of a right nullspace basis of $\varphi(A)$ over \mathbb{Z}_p .
3. Let $N = AN' \in \mathbb{Z}_{p^e}^{\ell \times k}$. This involves k matrix vector products with A . Note that N is divisible by p .
4. Compute the Smith normal form of N over \mathbb{Z}_{p^e} by Gaussian elimination:

$$\text{diag}(\underbrace{p, \dots, p}_{r_1}, \underbrace{p^2, \dots, p^2}_{r_2}, \dots, \underbrace{p^{e-1}, \dots, p^{e-1}}_{r_{e-1}}, \underbrace{0, \dots, 0}_{r_e}).$$

5. Return r_0, \dots, r_{e-1} .

We will analyze the algorithm holding e and p constant. Considering them as parameters would introduce a factor of $\mathcal{O}(e \log(p))$. Let the cost of matrix-vector product by B is $\mathcal{O}(\mu)$. Since we are holding e and p constant, this is the same for application to vectors in \mathbb{Z}_p^n and in $\mathbb{Z}_{p^e}^n$.

Step 0: Toeplitz matrices may be applied to vectors via polynomial multiplication, so the cost of the black-box for A is $\mathcal{O}(M(n) + \mu)$. $M(n)$ is $\mathcal{O}(n)$

and we will assume $\mu \geq n$, so the black-box cost of matrix-vector product by A is $\mathcal{O}(\mu)$.

Step 1: The rank over \mathbb{Z}_p can be done by a black-box method in $\mathcal{O}((\ell\mu) \log(\xi))$ to achieve probability of error less than $1/\xi$ (Wiedemann, 1986). Memory requirement is $\mathcal{O}(1)$ vectors in \mathbb{Z}_p^ℓ .

Step 2: Let $k = \ell - r_0$ denote the nullity of A modulo p . By black-box methods, k random samples of the nullspace will yield a nullspace basis N' . Oversampling can be done and column echelon form computation used to reduce to a basis of k columns if need be. The cost is $\mathcal{O}(k(\ell\mu))$. Space is $\mathcal{O}(k\ell)$. For instance see Chen et al. (2002).

Step 3: The cost of applying A to N' is $\mathcal{O}(k\mu)$.

Step 4: Any nullspace for S over \mathbb{Z}_p is of the form EW' , where E is the last $\ell - r_0$ columns of the identity and W' is $k \times k$ unimodular. Then $0 = AN' = USVN = USEW'$ modulo p , for some unimodular W' . This lifts to a factorization $AN = USEW$ modulo p^e with U, W unimodular. Thus AN has the Smith form SE . The local Smith form of this dense matrix can be computed by elimination. An algorithm running in $\mathcal{O}(\ell k^{\omega-1})$ is in Storjohann (2000).

Theorem 16. *Algorithm 2 is a correct Monte Carlo algorithm for computing Smith normal form over \mathbb{Z}_{p^e} . The time complexity is $\mathcal{O}(\ell k(k^{\omega-2} + \mu))$, where k is the number of nontrivial (neither 0 nor 1) invariant factors, ℓ is the reduced dimension (which can be the rank), and μ is the cost of matrix vector product by B . Note that the time complexity is $\mathcal{O}(\ell k\mu)$ under the very modest assumption that $k^{\omega-2} < \mu$. The memory requirement is $\mathcal{O}(k\ell)$.*

3.5 A conjecture about p -adic carries

While attempting a general approach for the \mathbb{Z}_{p^e} case without using dense elimination and while using only rank computations over the field \mathbb{Z}_p , we discovered an interesting pattern about ranks of matrices over \mathbb{Z}_{p^e} which could be of independent interest. To put the discovered conjecture in a proper context, we first introduce the attempt which lead to discovering it, then we present the conjecture which discourages this approach.

Consider an element $\alpha \in \mathbb{Z}_{p^e}$ written in terms of its *unique p -adic expansion* as $\alpha = \alpha_0 + \alpha_1 p + \dots + \alpha_{e-1} p^{e-1}$ where $\alpha_i \in \mathbb{Z}_p$ for $0 \leq i \leq e-1$. This representation extends naturally to matrices over \mathbb{Z}_{p^e} , i.e., $A = A_0 + pA_1 + \dots + p^{e-1}A_{e-1}$, where $A_i \in \mathbb{Z}_p^{n \times n}$. For clarity of presentation and limited space,

we only confine the discussion to the case $e = 2$. Expanding $A = USV$, we get $(A_0 + pA_1) = (U_0 + pU_1)(S_0 + pS_1)(V_0 + pV_1)$, where $A_0 = U_0S_0V_0 \bmod p$ and

$$A_1 = U_1S_0V_0 + U_0S_0V_1 + U_0S_1V_0 \quad (6)$$

where $\text{rank}(S_0) = r_0$, $\text{rank}(S_1) = r_1$, which are the multiplicities of 1's and p 's in the Smith form diagonal, respectively. Furthermore, with appropriate preconditioning*, $\text{rank}(A_0)$ is proportional to r_0 , and $\text{rank}(A_1)$ is proportional to $r_0 + 2r_1$. Hence, this formulation leads to a belief that we can easily *isolate* A_0, A_1 , compute their ranks over \mathbb{Z}_p , and discover multiplicities of the invariant factors. However, by closer inspection, equation 6 is in fact:

$$A_1 = U_1S_0V_0 + U_0S_0V_1 + U_0S_1V_0 + \overbrace{U_0S_0V_0 \text{ quo } p}^{\text{carry}},$$

where the extra term, $(U_0S_0V_0 \text{ quo } p)$, is introduced by the fact that operations in \mathbb{Z}_{p^e} exhibit carries. As a simple example $5 \cdot 5$ over \mathbb{Z}_{3^4} is $(2+p)(2+p) = 1 + 2p + 2p^2$, where $2p^2$ is a carry term. These carries contribute to the overall ranks of matrix expressions, and present a challenge in computing the Smith form. We hoped to reasonably bound the ranks of these carries, so a working algorithm could be developed which reduces Smith form computation to efficient rank computations over \mathbb{Z}_p . This approach would be superior to algorithms presented in previous section and literature since it preserves sparsity. It is worth noting that in the polynomial case, this approach yields a working algorithm. The following conjecture illustrates why the p -adic case is more difficult to resolve in this way.

Conjecture 17. Assume p is a prime, $U, S, V \in \mathbb{Z}_p^{n \times n}$ such that U, V are invertible, $S = \text{diag}(1, \dots, 1, 0, \dots, 0)$ and $\text{rank}(S) = r$. Let $M = USV = M_0 + M_1p + \dots + M_sp^s$, where $s = \mathcal{O}(\log_p r)$ and $M_i \in \mathbb{Z}_p^{n \times n}$. We conjecture that when $p = 2$,

$$\text{rank}(M_i) \leq \binom{r}{2^i},$$

and when $p = 2k + 1$, we conjecture that

$$\text{rank}(M_1) \leq \sum_{i=0}^k \binom{r+2i}{2i+1} + \binom{r+2k-1}{2k} - 2r.$$

*Details are outside the scope of presenting this conjecture.

Furthermore, in the generic case where U, V are uniformly chosen at random in \mathbb{Z}_p , and n is arbitrarily large, the ranks are equal to bounds above.

This conjecture shows that a large dimension product of matrices with entries in \mathbb{Z}_p of small rank can still have very large, but not full, rank carry matrices. These carries will impact many digits in the expanded product. The evidence for the conjecture is experimental[†]. We developed the formulas above by reverse engineering sequences of computed ranks resulting from experiments with different primes and matrices of different dimensions.

In the generic case of equality, the conjecture could be used to generate an algorithm for small primes, e.g. when $p = 2$. However, without further refinements, this approach yields an exponential time algorithm which is prohibitive.

4 Conclusion

We have given efficient algorithms for computing Smith Normal Form over two local rings: $\mathbb{F}[z]/(f^e)$ and $\mathbb{Z}/p^e\mathbb{Z}$. These are useful components for SNF algorithms over $\mathbb{F}[z]$ and \mathbb{Z} , respectively. These algorithms are efficient in the black-box model, which means that they are well suited to sparse and structured matrices. The integer algorithm is output sensitive. Its memory and time usage grows in proportion to the number of nontrivial invariant factors. A memory efficient algorithm without that restriction has not been found. In addition, we gave a conjecture about the rank of carries resulting from multiplying single-digit p -adic matrices.

References

- E. Babson, A. Björner, S. Linusson, J. Shareshian, and V. Welker. Complexes of not i -connected graphs. *Topology*, 38(2):271–299, 1999.
- A. Björner and V. Welker. Complexes of directed graphs. *SIAM Journal on Discrete Mathematics*, 12(4):413–424, November 1999.
- D. B. Chandler. Private communication, 2011. Symplectic matrix examples generated by Peter Vandendriessche.

[†]Code is available at: <http://cs.uwaterloo.ca/~mwg/smithpe/>

- D. B. Chandler, P. Sin, and Q. Xiang. Incidence modules for symplectic spaces in characteristic two. *J. Algebra*, 323:3157–3181, 2010.
- L. Chen, W. Eberly, E. Kaltofen, B. D. Saunders, W. J. Turner, and G. Villard. Efficient matrix preconditioners for black box linear algebra. *Linear Algebra and its Applications*, 343–344:119–146, 2002.
- Z. Chen and A. Storjohann. Lattice compression of integer matrices. *Journal of Symbolic Computation*, 2005. Accepted for publication.
- T. J. Chou and G. E. Collins. Algorithms for the solution of systems of linear diophantine equations. *SIAM Journal on Computing*, 11(4):687–708, 1982.
- D. Coppersmith and S. Winograd. Matrix multiplication via arithmetic progressions. *Journal of Symbolic Computation*, 9(3):251–280, 1990.
- J.-G. Dumas, B. D. Saunders, and G. Villard. Integer Smith form via the valence: Experience with large sparse matrices from homology. In *Proc. 2000 Internat. Symp. Symbolic Algebraic Comput. ISSAC’00*, pages 95–105. ACM Press, 2000.
- J.-G. Dumas, B. D. Saunders, and G. Villard. On efficient sparse integer matrix Smith normal form computations. *Journal of Symbolic Computation*, 32:71–99, 2001.
- J.-G. Dumas, F. Heckenbach, B. D. Saunders, and V. Welker. Computing simplicial homology based on efficient Smith normal form algorithms. In *Algebra, Geometry, and Software Systems*, pages 177–206. Springer, 2003.
- J.-G. Dumas, C. Pernet, and B. D. Saunders. On finding multiplicities of characteristic polynomial factors of black-box matrices. In *Proc. International Symposium on Symbolic and Algebraic Computation (ISSAC’09)*, pages 135–142. ACM, 2009.
- W. Eberly. Reliable Krylov-based algorithms for matrix null space and rank. In *Proc. ISSAC 2004*, pages 127–134, 2004.
- W. Eberly, M. Giesbrecht, and G. Villard. On computing the determinant and smith form of an integer matrix. In *Proc. 41st Annual IEEE Symposium on Foundations of Computer Science (FOCS’2000)*, pages 675–687, 2000.
- W. Eberly, M. Giesbrecht, P. Giorgi, A. Storjohann, and G. Villard. Faster inversion and other black box matrix computations using efficient block projections. In *Proc. International Symposium on Symbolic and Algebraic Computation (ISSAC’07)*, pages 143–150, 2007.
- M. Giesbrecht. Fast computation of the Smith form of a sparse integer matrix. *Computational Complexity*, 10(1):41–69, 2001.
- J. L. Hafner and K. S. McCurley. A rigorous subexponential algorithm for com-

- putation of class groups. *Journal of the American Mathematical Society*, 2(4): 837–850, 1989.
- T. C. Hu. *Integer programming and network flows*. Addison-Wesley, Reading, Mass., 1969.
- T. Kailath. *Linear Systems*. Prentice-Hall, Englewood Cliffs, NJ, 1980.
- E. Kaltofen and B. D. Saunders. On Wiedemann’s method of solving sparse linear systems. In *Applied Algebra, Algebraic Algorithms and Error-Correcting Codes (AAECC ’91)*, volume 539 of *LNCS*, pages 29–38, 1991.
- E. Kaltofen, M.S. Krishnamoorthy, and B. David Saunders. Parallel algorithms for matrix normal forms. *Linear Algebra and its Applications*, 136:189–208, Jul 1990.
- I. Kaplansky. Elementary divisors and modules. *Transactions of the American Mathematical Society*, 66(2):464–491, 1949.
- B. McDonald. *Finite Rings with Identity*. Marcel Dekker, Inc., New York, 1974.
- B. McMillan. On systems of linear indeterminate equations and congruences. *Bell System Technical Journal*, 31:541–600, 1952.
- M. Newman. *Integral Matrices*. Academic Press, New York, NY, USA, 1972.
- B. D. Saunders and Z. Wan. Smith normal form of dense integer matrices fast algorithms into practice. In *Proc. 2004 Internat. Symp. Symbolic Algebraic Comput. ISSAC’04*, pages 274–281. ACM Press, 2004.
- B. D. Saunders and B. Youse. Large matrix, small rank. In *Proc. International Symposium on Symbolic and Algebraic Computation (ISSAC’09)*, pages 317–324, 2009.
- J. T. Schwartz. Fast probabilistic algorithms for verification of polynomial identities. *Journal of the ACM*, 27(4):701–717, October 1980.
- H. J. S. Smith. On systems of linear indeterminate equations and congruences. *Philosophical Transactions of the Royal Society of London*, 151:293–326, 1861.
- A. Storjohann. *Algorithms for matrix canonical forms*. PhD thesis, Swiss Federal Institute of Technology – ETH, Zürich, 2000.
- A. Storjohann. High-order lifting and integrality certification. *Journal of Symbolic Computation*, 36(3–4):613–648, 2003.
- W.D. Wallis, A. P. Street, and J. S. Wallis. *Combinatorics: room squares, sum-free sets, Hadamard matrices*. Lecture notes in mathematics. Springer, Berlin, 1972.
- D. Wiedemann. Solving sparse linear equations over finite fields. *IEEE Transactions on Information Theory*, IT-32:54–62, January 1986.

- J. Wilkening and J. Yu. A local construction of the Smith normal form of a matrix polynomial. *Journal of Symbolic Computation*, 46(1):1–22, January 2011.
- R. Zippel. Probabilistic algorithms for sparse polynomials. In *Symbolic and Algebraic Computation*, volume 72 of *LNCS*, pages 216–226. Springer Berlin, 1979.